

CLAIMS

What is claimed is:

- 5 1. A method of transferring data from a sender to a receiver in a communication network, the method comprising:
 - establishing a connection between the sender and the receiver;
 - measuring round trip times of data packets sent from the sender to the receiver;
 - determining a congestion window parameter that specifies a maximum number of
 - 10 unacknowledged data packets that may be sent to the receiver; and
 - transmitting additional data packets to the receiver in response to expiration of a transmit timer, the period of the transmit timer based on the round trip time measurements and the congestion window parameter.
- 15 2. The method of claim 1, wherein the step of measuring the round trip times comprises measuring the time between transmission of a data packet and receipt of a corresponding acknowledgement.
- 20 3. The method of claim 1, wherein the step of determining comprises adjusting the congestion window parameter in response to receipt of acknowledgements from the receiver.
- 25 4. The method of claim 3, wherein the step of adjusting comprises incrementing the congestion window parameter in response to receipt of each acknowledgement.
- 30 5. The method of claim 3, wherein the step of adjusting comprises incrementing the congestion window parameter in response to receipt of a number of acknowledgments equal to the size of the then-current congestion window parameter.
6. The method of claim 3, wherein the step of adjusting comprising increasing the congestion window parameter by an amount equal to a number of buffers effectively freed by each acknowledgement.

7. The method of claim 1, wherein the period of the timer is based on a smoothed round trip time and a smoothed congestion window.

8. The method of claim 7, wherein the period of the timer is based on a ratio of the smoothed round trip time and the smoothed congestion window.

9. The method of claim 7, wherein the smoothed round trip time is based on an average and mean deviation of a plurality of round trip time measurements.

10. The method of claim 7, wherein the smooth round trip time is based on a weighted sum of a previous smoothed round trip time and a then-current round trip time measurement.

11. The method of claim 7, wherein the smoothed congestion window comprises a weighted sum of a previous smoothed congestion window and a then-current congestion window parameter.

12. The method of claim 1, wherein the step of transmitting comprises multiplexing new and retransmitted data packets.

13. The method of claim 12, wherein the step of multiplexing comprises determining whether to transmit either the new or the retransmitted data packets based on expiration of a timer.

14. The method of claim 12, wherein the step of multiplexing comprises determining whether to transmit either the new or the retransmitted data packets based on a number of duplicate acknowledgments exceeding a predetermined threshold.

15. A method for responding to packet loss, the method comprising:
setting a short timer and a long timer in response to transmission of a data packet,
the long timer having a longer period than the short timer;

if a number of unacknowledged data packets does not exceed a predetermined threshold, retransmitting the data packet at a first transmission rate in response to each expiration of the short timer;

resetting the short timer and the long timer in response to each retransmission of the data packet; and

reducing the first transmission rate, if the long timer expires before receipt of a corresponding acknowledgement for the data packet.

16. The method of claim 15, wherein the period of the short timer is based on a smoothed round trip time of a plurality of round trip time measurements and a maximum smoothed round trip time determined over a predetermined time period.

17. The method of claim 16, wherein the period of the short timer is based on a minimum of the smoothed round trip time and the maximum smoothed round trip time.

18. The method of claim 15, wherein the period of the long timer is based on a smoothed round trip time of a plurality of round trip time measurements and a maximum variation of the plurality of round trip time measurements.

19. The method of claim 18, wherein the period of the long timer is based on a maximum of the smoothed round trip time and the maximum variation of the plurality of round trip time measurements.

20. The method of claim 15, further comprising incrementing the predetermined threshold to allow retransmission of the data packet, if the number of unacknowledged data packets exceeds the predetermined threshold and the data packet is retransmitted for the first time.

21. The method of claim 15, further comprising using a transmit timer to control transmission and retransmission of data packets, the period of the transmit timer determining the transmission rate of the data packets.

22. The method of claim 21, wherein the step of retransmitting the data packet at the first transmission rate is performed in response to expiration of both the short timer and the transmit timer.

5 23. The method of claim 21, wherein the period of the transmit timer is based on a ratio of a round trip time of the data packets and a congestion window parameter that specifies a maximum number of unacknowledged data packets.

10 24. The method of claim 23, wherein the step of reducing comprises resetting the congestion window parameter to an initial value to increase the period of the transmit timer.

15 25. The method of claim 15, wherein the step of reducing is performed such that the first transmission rate is not reduced until after at least a second retransmission of the data packet.

20 26. A method for adjusting a rate of data transmission within a communication network, the method comprising:

transmitting data packets at a transmission rate based on a size of a congestion window;

measuring a round trip time for each of the data packets;

increasing the size of the congestion window at a first rate until a smoothed round trip time of the round trip time measurements exceeds a predetermined threshold; and

25 increasing the size of the congestion window at a second rate after the smoothed round trip time of the round trip time measurements exceeds the predetermined threshold, the second rate being smaller than the first rate.

30 27. The method of claim 26, wherein the transmission rate is based on a smoothed round trip time and a smoothed congestion window.

28. The method of claim 27, wherein the transmission rate is based on a ratio of the smoothed round trip time and the smoothed congestion window.

29. The method of claim 27, wherein the smoothed congestion window comprises a weighted sum of a previous smoothed congestion window and a then-current congestion window.

5 30. The method of claim 27, wherein the smoothed round trip time is based on an average and mean deviation of a plurality of the round trip time measurements.

10 31. The method to claim 26, wherein the step of measuring the round trip times comprises measuring the time between transmission of a data packet and receipt of a corresponding acknowledgement.

32. The method of claim 26, wherein the predetermined threshold comprises a minimum smoothed round trip time over a predetermined time period.

15 33. The method of claim 26, wherein the step of increasing the size of the congestion window at the first rate comprises incrementing the size of the congestion window for each received acknowledgement.

20 34. The method of claim 26, wherein the step of increasing the size of the congestion window at the first rate comprises increasing the size of the congestion window by a number of buffers effectively freed by an acknowledgement.

25 35. The method of claim 26, wherein the step of increasing the size of the congestion window at the second rate comprises incrementing the size of the congestion window in response to receiving a number of acknowledgements equal to the then-current size of the congestion window.

36. A method for avoiding congestion within a communication network, the method comprising:

30 transmitting data packets at a transmission rate determined based on a size of a congestion window;

measuring a round trip time for each of the data packets;

determining whether a change in round trip time measurements exceeds a threshold value;

increasing the transmission rate by increasing the size of the congestion window in response to the change in round trip time measurements not exceeding the threshold value;
and

reducing the transmission rate by reducing the size of the congestion window in response to the change in round trip time measurements exceeding the threshold value.

37. The method of claim 36, wherein the transmission rate is based on a smoothed round trip time and a smoothed congestion window.

38. The method of claim 37, wherein the transmission rate is based on a ratio of the smoothed round trip time and the smoothed congestion window.

39. The method to claim 36, wherein the step of measuring the round trip times comprises measuring the time between transmission of a data packet and receipt of a corresponding acknowledgement.

40. The method of claim 36, wherein the step of increasing the size of the congestion window comprises incrementing the size of the congestion window for each received acknowledgement.

41. The method of claim 36, wherein the step of increasing the size of the congestion window comprises increasing the size of the congestion window by a number of buffers effectively freed by an acknowledgement.

42. The method of claim 36, wherein the step of determining is performed in response to receipt of each acknowledgement associated with the data packets, and wherein the step of reducing the transmission rate is performed in response the change in the round trip time measurements exceeding the threshold value for at least two of the acknowledgements.

43. The method of claim 36, further comprising incrementing the size of the congestion window in response to each reduction in the size of the congestion window.

44. A method for managing a plurality of connections between a sender and a receiver, the method comprising:

determining a transmission rate for the plurality of connections;
calculating a weight for each of the plurality of connections based on a respective congestion window parameter associated with each connection; and

transmitting a data packet associated with a selected one of the plurality of connections having the highest weight to provide a weight-based allocation of the transmission rate among the plurality of connections.

45. The method of claim 44, further comprising initializing a congestion window parameter for a new connection between the sender and the receiver based on a minimum smoothed round trip time for the plurality of connections.

46. The method of claim 44, further comprising initializing a congestion window parameter for a new connection between the sender and the receiver based on an identification of the receiver.

47. The method of claim 44, further comprising adjusting the congestion window parameter associated with the selected connection in response to receipt of an acknowledgement for the transmitted data packet.

48. The method of claim 47, wherein the step of adjusting comprises incrementing the congestion window parameter in response to receipt of the acknowledgement.

49. The method of claim 47, wherein the step of adjusting comprises reducing the congesting window parameter in response to the sum of the change in congestion window parameters for the plurality of connections exceeding a predetermined threshold.

50. A system for transferring data from a sender to a receiver in a communication network, the system comprising:

a processor; and

a memory unit operably coupled to the processor, the memory unit storing instructions which when executed by the processor cause the processor to operate so as to:

establish a connection between the sender and the receiver;

measure round trip times of data packets sent from the sender to the receiver;

determine a congestion window parameter that specifies a maximum number of unacknowledged data packets that may be sent to the receiver; and

transmit additional data packets to the receiver in response to expiration of a transmit timer, the period of the transmit timer based on the round trip time measurements and the congestion window parameter.

51. A system for responding to packet loss, the system comprising:

a processor; and

a memory unit operably coupled to the processor, the memory unit storing instructions which when executed by the processor cause the processor to operate so as to:

set a short timer and a long timer in response to transmission of a data packet, the long timer having a longer period than the short timer;

if a number of unacknowledged data packets does not exceed a predetermined threshold, retransmit the data packet at a first transmission rate in response to each expiration of the short timer;

reset the short timer and the long timer in response to each retransmission of the data packet; and

reduce the first transmission rate, if the long timer expires before receipt of a corresponding acknowledgement for the data packet.

52. A system for adjusting a rate of data transmission within a communication network, the system comprising:

a processor; and

a memory unit operably coupled to the processor, the memory unit storing instructions which when executed by the processor cause the processor to operate so as to:

transmit data packets at a transmission rate based on a size of a congestion window;

measure a round trip time for each of the data packets;

increase the size of the congestion window at a first rate until a smoothed round trip time of the round trip time measurements exceeds a predetermined threshold; and
increase the size of the congestion window at a second rate after the smoothed round trip time of the round trip time measurements exceeds the predetermined threshold, the second rate being smaller than the first rate.

53. A system for avoiding congestion within a communication network, the system comprising:

a processor; and

a memory unit operably coupled to the processor, the memory unit storing instructions which when executed by the processor cause the processor to operate so as to:

transmit data packets at a transmission rate determined based on a size of a congestion window;

measure a round trip time for each of the data packets;

determine whether a change in round trip time measurements exceeds a threshold value;

increase the transmission rate by increasing the size of the congestion window in response to the change in round trip time measurements not exceeding the threshold value; and

reduce the transmission rate by reducing the size of the congestion window in response to the change in round trip time measurements exceeding the threshold value.

54. A system for managing a plurality of connections between a sender and a receiver, the system comprising:

a processor; and

a memory unit operably coupled to the processor, the memory unit storing instructions which when executed by the processor cause the processor to operate so as to:

determine a transmission rate for the plurality of connections;

calculate a weight for each of the plurality of connections based on a respective congestion window parameter associated with each connection; and

transmit a data packet associated with a selected one of the plurality of connections having the highest weight to provide a weight-based allocation of the transmission rate among the plurality of connections.

transmit a data packet associated with a selected one of the plurality of connections having the highest weight to provide a weight-based allocation of the transmission rate among the plurality of connections.